

Robotics Research Technical Report

Arrangements of Curves in the Plane — Topology, Combinatorics, and Algorithms

by


Herbert Edelsbrunner, Leonidas Guibas,
Janos Pach, Richard Pollack, Raimund Seidel,
Micha Sharir

Technical Report No. 366
Robotics Report No. 151
April, 1988

NYU COMPSCI TR-366
Edelsbrunner, Herbert
Arrangements of curves in
the plane - topology,
combinatorics,... c.2

New York University
Institute of Mathematical Sciences

Computer Science Division
251 Mercer Street New York, N.Y. 10012



**Arrangements of Curves in the Plane —
Topology, Combinatorics, and Algorithms**

by

Herbert Edelsbrunner, Leonidas Guibas,
Janos Pach, Richard Pollack, Raimund Seidel,
Micha Sharir

Technical Report No. 366
Robotics Report No. 151
April, 1988

New York University
Dept. of Computer Science
Courant Institute of Mathematical Sciences
251 Mercer Street
New York, New York 10012

Work on this paper by the first author has been supported by the National Science Foundation under grant CCR-8714565. Work by the third and sixth authors has been supported by Office of Naval Research Grant N00014-82-K-0381, by National Science Foundation Grant No. NSF-DCR-83-20085, by grants from the Digital Equipment Corporation, and the IBM Corporation. Work by the sixth author has also been supported by a research grant from the NCRD — the Israeli National Council for Research and Development. Work by the fourth author has been supported by National Science Foundation Grant DMS-8501947.

Arrangements of Curves in the Plane – Topology, Combinatorics, and Algorithms

Herbert Edelsbrunner⁽¹⁾, Leonidas Guibas^(2,3), Janos Pach^(4,5),
Richard Pollack⁽⁴⁾, Raimund Seidel⁽⁶⁾ and Micha Sharir^(4,7)

- (1) Computer Science Department, University of Illinois at Urbana
- (2) DEC Systems Research Center
- (3) Computer Science Department, Stanford University
- (4) Courant Institute of Mathematical Sciences, New York University
- (5) Mathematical Institute of the Hungarian Academy of Sciences
- (6) Computer Science Department, University of California at Berkeley
- (7) School of Mathematical Sciences, Tel Aviv University

ABSTRACT

Arrangements of curves in the plane are fundamental to many problems in computational and combinatorial geometry (e.g. motion planning, algebraic cell decomposition, etc.). In this paper we study various topological and combinatorial properties of such arrangements under some mild assumptions on the shape of the curves, and develop basic tools for the construction, manipulation, and analysis of these arrangements. Our main results include a generalization of the zone theorem of [EOS], [CGL] to arrangements of curves (in which we show that the combinatorial complexity of the zone of a curve is nearly linear in the number of curves), and an application of that theorem to obtain a nearly quadratic incremental algorithm for the construction of such arrangements.

1. Introduction

A *Jordan arc* γ is the image of a continuous one-to-one mapping from the interval $[0,1]$ to the plane (together with a point at infinity). If the removal of γ decomposes the plane into two connected components then γ is called a *Jordan curve*. If, in addition, γ is bounded then it is necessarily the boundary of a bounded and simply connected set in the plane; in this case γ is said to be *closed*.

Work on this paper by the first author has been supported by the National Science Foundation under grant CCR-8714565. Work by the third and sixth authors has been supported by Office of Naval Research Grant N00014-82-K-0381, by National Science Foundation Grant No. NSF-DCR-83-20085, by grants from the Digital Equipment Corporation, and the IBM Corporation. Work by the sixth author has also been supported by a research grant from the NCRD – the Israeli National Council for Research and Development. Work by the fourth author has been supported by National Science Foundation Grant DMS-8501947.

Let $\Gamma = \{\gamma_1, \gamma_2, \dots, \gamma_n\}$ be a collection of n closed or unbounded Jordan curves in the plane, such that each pair of these curves intersect transversally in at most s points, for some fixed integer s . The *arrangement* $A = A(\Gamma)$ of Γ is the planar map induced by these curves; this is a subdivision of the plane whose *vertices* are the intersection points of the curves γ_i , whose *edges* are the connected components of these curves minus the vertices, and whose *faces* are maximal connected regions of the complement of the union of these curves. We will assume in what follows that the arrangement A is *simple*, meaning that no point is common to three or more curves (and, as already assumed, no two curves in Γ have a point in common where they do not cross). Since any two curves in Γ intersect in at most s points, it follows that the number of vertices in A is at most $sn(n-1)/2$, and an easy application of Euler's formula allows us to conclude that the number of edges and faces in A is also $O(n^2)$.

We will also consider the case where the curves γ_i are bounded Jordan arcs. In this case (as in the case of unbounded Jordan curves where we do not count intersections at the point at infinity), the maximum number s of intersections between any pair of these curves can be odd. The arrangement A of such arcs is defined in the same way as above, except that the endpoints of the γ_i 's are also taken to be vertices of A , and that simplicity of A now also requires that no endpoint of one arc γ_i lies on another arc γ_j .

Arrangements of curves arise naturally in many problems in computational geometry. For example, Chazelle and Lee [CL] consider the following problem. Given n points x_1, x_2, \dots, x_n in the plane with associated weights w_1, w_2, \dots, w_n , and a fixed radius r , find a placement of a circle of radius r that maximizes the sum of the weights of the points x_i lying within the circle. It is shown in [CL] that this problem can be reduced to the problem of calculating the arrangement of n circles of radius r centered at each of the given points, and then searching through the faces of this arrangement to find an optimal placement. We can generalize this problem in several ways. For example, let C_1, C_2, \dots, C_n be n closed convex and disjoint sets in the plane, and let B be another convex set. Allow B to translate in the plane, and for each placement of B we want to count how many objects C_i it intersects, or sum up certain weights associated with them, etc. Using standard techniques (as in [KLPS]), we can form the Minkowsky (vector) differences $K_i = C_i - B_0$, where B_0 is some standard placement of B in which some fixed reference point O of B lies at the origin. Let γ_i be the boundary of K_i , for $i=1, 2, \dots, n$. It is easily checked that, for each face f of the arrangement $A = A(\{\gamma_1, \gamma_2, \dots, \gamma_n\})$, all placements of B with the reference point O lying in f intersect the same subset of the objects C_i . Thus calculation of A will provide a solution to the problems stated above. Note that in this case $\gamma_i \neq \gamma_j$ intersect in at most $s=2$ points if we assume general position (see [KLPS] for a more precise statement of this condition and for a proof of this property).

Many other problems can be reduced to the analysis of an arrangement of curves. This is the case, for example, when the objects under consideration can be parametrized as points in the plane, and certain properties of such an object vary discontinuously as it crosses certain "critical curves". The arrangement of these curves partitions the plane into "non-critical regions", and construction of these regions is often required to obtain a discrete combinatorial representation of all possible problem states. Such examples, involving motion planning problems, can be found in [SS1], [KO], [MO], [GSS].

A special case of arrangements which has been studied extensively in the past is that of lines. An important property of such arrangements is the so-called "Zone Theorem" (see [CGL], [EOS], [Ed]) which states the following.

Theorem. Let A be an arrangement of n lines l_1, l_2, \dots, l_n , and let l be another line. Then the total number of edges bounding the faces of A that intersect l is $O(n)$.

We refer to the collection of all these edges as the *zone* of l in A .

One useful application of this theorem is that it facilitates the construction of the arrangement A_{n+1} of $n+1$ lines $l_1, l_2, \dots, l_n, l_{n+1}$ from the arrangement A_n of the first n lines in linear time as follows. Assume without loss of generality that $l = l_{n+1}$ is the x -axis. First find the leftmost unbounded face of A_n crossed by l . Next process the faces of A_n crossed by l from left to right. At each such face f find the rightmost point of $l \cap f$; this will determine the next face f' of A_n crossed by l , and the process is then repeated for f' . The crossing points of l with the boundaries of the faces in A_n are found by traversing all edges in the zone of l ; the number of such edges is $O(n)$ by the Zone Theorem. For each of these faces f , the algorithm also splits f into two new faces in A_{n+1} , and updates (also in linear time) the planar map representation of the arrangement. The resulting sequence of incremental updates yields an overall optimal $O(n^2)$ algorithm for the calculation of arrangements of n lines.

The zone theorem has found other applications, some of which are described in [Ed].

Our goal is to extend the study of planar arrangements to allow more general curves. In particular, we want to extend the zone theorem to such arrangements and explore its algorithmic consequences. Of course, in this more general setting, we encounter several new technical difficulties which make analysis and calculation of these arrangements somewhat more complicated. First of all, the bound given in the zone theorem for lines may be incorrect for certain collections of (rather simple) curves, such as line segments or circles, and a modified estimate of the overall complexity of the arrangement faces crossed by a new curve is required. To expand upon this issue, let Γ_0 be an (infinite) collection of Jordan curves or arcs with the property that any two non-overlapping curves in Γ_0 (that is, curves whose intersection does not contain any arc) intersect in at most s points. For each finite subcollection $\Gamma \subset \Gamma_0$ of

non-overlapping curves that give rise to a simple arrangement $A(\Gamma)$, and for each curve $\gamma \in \Gamma_0 - \Gamma$ which does not overlap any $\gamma \in \Gamma$ and such that $A(\Gamma \cup \{\gamma\})$ is also simple, let $\mu(\Gamma, \gamma)$ denote the total number of edges in the faces of $A(\Gamma)$ crossed by γ (again we refer to the collection of all these edges as the *zone* of γ in $A(\Gamma)$). Let

$$\sigma_n(\Gamma_0) = \max\{\mu(\Gamma, \gamma)\},$$

where the maximum is taken over all n -element subcollections $\Gamma \subset \Gamma_0$ and curves $\gamma \in \Gamma_0$ satisfying the above conditions.

Our main result is

Theorem 1. $\sigma_n(\Gamma_0) = O(\lambda_{s+2}(n))$.

Here $\lambda_s(n)$ is the maximum length of a so-called (n, s) Davenport–Schinzel sequence (or (n, s) -sequence) defined as follows. An (n, s) -sequence is a linear sequence composed of n different elements so that no two consecutive elements are the same and there is no (not necessarily contiguous) subsequence of length $s+2$ of the form $\mu\nu\mu\nu\dots$, for $\mu \neq \nu$. We refer to [HS], [Sh1], [Sh2], [ASS] for more details concerning these sequences. The following is known about the length of Davenport–Schinzel sequences.

$$\lambda_1(n) = n \text{ and } \lambda_2(n) = 2n - 1 \text{ (trivial).}$$

$$\lambda_3(n) = \Theta(n\alpha(n)), \text{ where } \alpha(n) \text{ is the functional inverse of Ackermann's function, and thus grows extremely slowly [HS].}$$

$$\lambda_4(n) = \Theta(n \cdot 2^{\alpha(n)}) \text{ [ASS].}$$

$$\lambda_{2s}(n) = n \cdot 2^{\Theta(\alpha(n)^{s-1})} \text{ for } s > 2 \text{ [ASS].}$$

$$\lambda_{2s+1}(n) = n \cdot \alpha(n)^{O(\alpha(n)^{s-1})} \text{ for } s \geq 2 \text{ [ASS].}$$

Thus $\lambda_s(n)$ is almost linear in n for any fixed s (and superlinear for $s \geq 3$).

For some example applications of our main result, let \mathbf{L} denote the collection of all lines in the plane, \mathbf{S} denote the collection of all line segments, and \mathbf{C} denote the collection of all circles. For \mathbf{L} , an appropriate modification of our proof technique yields an improved bound

$$\sigma_n(\mathbf{L}) = O(\lambda_2(n)) = O(n),$$

which matches the bound in [EOS], [CGL]. For \mathbf{S} our result specializes to

$$\sigma_n(\mathbf{S}) = O(\lambda_3(n)) = O(n\alpha(n)),$$

and for \mathbf{C} we obtain

$$\sigma_n(\mathbf{C}) = O(\lambda_4(n)) = O(n \cdot 2^{\alpha(n)}).$$

Moreover, recent results by Wiernik and Sharir [WS] and by Shor [Sho] imply a lower bound of $\Omega(n\alpha(n))$ on both $\sigma_n(\mathbf{S})$ and $\sigma_n(\mathbf{C})$.

When we try to calculate an arrangement of more general curves, several additional difficulties arise. One difficulty is that the collection of edges in such arrangements need not be connected. Another is that individual faces in such arrangements can have a fairly complex structure. The incremental technique of [EOS], [CGL] for arrangements of lines is based on the property that each face f in such an arrangement is convex; thus any line intersects f in at most two points, and these points can be easily computed in time proportional to the number of edges of f . In arrangements of more general curves, a face f may have a disconnected boundary, and every single connected component of its boundary may intersect the additional curve γ in many points. Even though all these points can be calculated in time proportional to the number of edges of f (assuming that “primitive” operations, such as calculating the intersection points of a given pair of curves, require constant time), appropriate sorting of these points along γ , in constant time per point, is no longer a simple task to accomplish.

All these problems can be dealt with, even when using a naive representation, where the faces of the current arrangement are maintained by the collections of their boundary components, each stored as a circular list, so that intersections of a new curve γ with these boundaries are found by simply traversing these lists. However, this approach requires several elaborate and complicated techniques (including Jordan sorting [HMRT] of the intersections along γ) which make implementation of this technique impractical.

Instead, we propose to maintain each face of the current arrangement using a *vertical cell decomposition* (which splits each face f into trapezoidal-like subcells) and then update this cell decomposition as a new curve is being added. This leads to a rather simple algorithm, whose efficient performance depends on our zone theorem. For this simpler method to apply, some further (though quite natural) restrictions have to be imposed on the shape of the curves in Γ_0 . The resulting algorithm runs in time $O(n\lambda_{s+2}(n))$, and is thus close to optimal in the worst case. Incidentally, our algorithm is quite similar to an independently discovered algorithm by Mulmuley [M] which constructs the arrangement of line segments. A main difference between our similar approaches is, however, that he gives a randomized time bound while our work is geared towards the worst case that can happen.

Our analysis in section 2 is based on several topological and combinatorial tools which are presented here. One tool of interest is Theorem 2 which provides a nearly linear upper bound on the complexity of a single face in an arrangement of curves.

We believe that the topological and combinatorial analysis given in section 2 will have many applications beyond the incremental construction of arrangements. For example, recent results in [CEGSW], [AS] use our generalized zone theorem in an analysis of the complexity of many cells in arrangements of curves, and in an analysis of the combinatorial complexity of a single component in an arrangement of triangles in three-dimensional space.

In the interest of brevity, throughout the paper we will use a somewhat informal language for describing operations on arrangements of curves.

2. Combinatorial Bounds for Zones

Let Γ_0 be a collection of Jordan curves or arcs with the properties stated in the introduction. Let $\Gamma = \{\gamma_1, \gamma_2, \dots, \gamma_n\}$ be a finite subcollection of Γ_0 , and let $\gamma \in \Gamma_0 - \Gamma$ be such that Γ and γ also satisfy the conditions stated in the introduction. That is, each pair of the curves in $\Gamma \cup \{\gamma\}$ intersect one another only transversally, in at most s points, and the arrangement $A(\Gamma \cup \{\gamma\})$ is simple.

Our goal is to obtain a sharp upper bound on the complexity of $\mu(\Gamma, \gamma)$, that is, the total number of edges of the faces of $A = A(\Gamma)$ that intersect γ . Since all our counting results will be asymptotical it does not make a difference whether or not we double-count an edge bounding two such faces.

We assume here that γ is a Jordan curve and thus separates the plane into two disjoint connected components, which we denote as K^+ and K^- . If γ is a bounded Jordan arc, we can “expand” γ into a closed Jordan curve by taking two disjoint curves lying very close to γ and connecting their endpoints; then we can apply our analysis to this curve. (Note that the resulting curve can now intersect any other curve in up to $2s$ points; this will, however, not change our analysis.)

Let Γ_γ be the collection of curves that intersect γ . Each $\gamma_i \in \Gamma_\gamma$ intersect γ in at most s points, so it is split by γ into at most $s+1$ connected pieces, each of which lies either in K^+ or in K^- .

Let Γ^+ (Γ^-) denote the collection of all curves in Γ (including the appropriate pieces of the curves in Γ_γ) that are fully contained in K^+ (K^-). Put $n^+ = |\Gamma^+|$, and $n^- = |\Gamma^-|$ (so $n^+ + n^-$ is proportional to n). In estimating $\mu(\Gamma, \gamma)$, it is clearly sufficient to bound only the number μ^+ of edges in the portions of the faces of A that are crossed by γ and lie in K^+ , and then use twice this bound as a bound for $\mu(\Gamma, \gamma)$.

Before starting our analysis, we first comment on the special case in which Γ_γ is empty and all curves in Γ are Jordan curves. For each $\gamma_i \in \Gamma^+$ let $K(\gamma_i)$ denote the connected component of $R^2 - \gamma_i$ that is disjoint from γ , and let $K = \bigcup_{\gamma_i \in \Gamma^+} K(\gamma_i)$.

Clearly γ is fully contained in a single connected component of the complement K^c of K . By the recent results of [SS2], the boundary of this component consists of at most $O(\lambda_s(n^+))$ edges. Hence $\mu(\Gamma^+, \gamma) = O(\lambda_s(n^+))$.

When this favorable case does not arise, we need to analyze the more general case in which some curves in Γ^+ may be Jordan arcs. In this case, Theorem 2 below implies the slightly weaker bound $\mu(\Gamma^+, \gamma) = O(\lambda_{s+2}(n^+))$. This theorem has originally been proven in [GSS] (and in [PSS] for the special case of line segments); for the sake of completeness we give below full details of the proof.

Theorem 2. Let $\Delta = \{\delta_1, \delta_2, \dots, \delta_m\}$ be a collection of m Jordan arcs, any two of which intersect in at most s points. Then the number of edges bounding a single face of $A(\Delta)$ is at most $O(\lambda_{s+2}(m))$.

Proof. Let f be the given face, and let C be a connected component of its boundary. It suffices to show that, if k arcs of Δ appear along C , then the number of edges of $A(\Delta)$ in C is $O(\lambda_{s+2}(k))$. Since $\lambda_{s+2}(k)$ is $\Omega(k)$ we may assume without loss of generality that all m arcs of Δ appear along C . For each δ_i let u_i, v_i be its endpoints. Let δ_i^+ (δ_i^-) be the directed arc δ_i oriented from u_i to v_i (from v_i to u_i).

Traverse C so that f lies to your left, and let $S = (s_1, s_2, \dots, s_l)$ be the circular sequence of oriented curves in Δ in the order in which they appear along C (if C is unbounded, then S is a linear rather than circular sequence). More precisely, if during our traversal of C we encounter a curve δ_i and follow it in the direction from u_i to v_i (resp. from v_i to u_i) then we add δ_i^+ (resp. δ_i^-) to S . As an example, if the endpoint u_i of δ_i is on C , then traversing C past u_i will add the pair of elements δ_i^-, δ_i^+ to S , and symmetrically for v_i (see Figure 1).

In what follows we will use the following notation. We denote the oriented arcs of Δ as $\xi_1, \xi_2, \dots, \xi_{2m}$. For each ξ_i we denote by $\overline{\xi_i}$ the non-oriented arc δ_i coinciding with ξ_i . For the purpose of the proof we will replace each arc δ_i by a closed Jordan curve δ_i^* that surrounds but does not intersect δ_i and whose points lie sufficiently close to δ_i . This will perturb the face f slightly, but, assuming the arrangement $A(\Delta)$ is simple, will not change the combinatorial structure of the boundary of f , and in particular of C . We can cut δ_i^* into two pieces at points close to the endpoints of δ_i so that one of the two pieces can be naturally identified with δ_i^+ and the other piece with δ_i^- .

We next need the following lemmas.

Lemma 3. The portions of each arc ξ_i appear in S in a circular order which is consistent with their order along the oriented ξ_i (that is, there exists a starting point in S - which depends on ξ_i - such that if we read S in circular order starting from that point, we encounter these portions in their order along ξ_i).

Proof. Let ζ, η be two portions of ξ_i that appear consecutively along C . Choose two points $x \in \zeta$ and $y \in \eta$ and connect them by the portion α of C traversed from x to y , and by another arc β inside $\overline{\xi_i}^*$. Clearly α and β do not intersect (except at their endpoints) and they are both contained in the complement of f . Thus their union $\alpha \cup \beta$ is a closed Jordan curve and f is either fully outside of fully inside $\alpha \cup \beta$ (see Figure 2). We claim that any point on ξ_i between ζ and η is contained in the side of $\alpha \cup \beta$ that does not contain f . Indeed, connect such a point z to x along an arc ρ that proceeds very near ξ_i outside $\overline{\xi_i}^*$ (see Figure 2). Clearly ρ and β are disjoint, and, deforming α slightly as necessary, we can assume that ρ intersects α transversally and exactly once, which implies our claim. This claim completes the proof of the

lemma. ■

For each directed arc ξ_i consider the linear sequence V_i of all appearances of ξ_i in S , arranged in the order they appear along ξ_i . Let f_i and l_i denote the index in S of the first and last element of V_i . Consider $S = (s_1, s_2, \dots, s_t)$ as a linear, rather than circular, sequence by cutting it at an arbitrary point (this step is not needed if C is unbounded). For each arc ξ_i , if $f_i > l_i$ we split ξ_i into two distinct arcs ξ_{i1} , ξ_{i2} , and replace all appearances of ξ_i in S between the places f_i and t (resp. between 1 and l_i) by ξ_{i1} (resp. ξ_{i2}). If we perform this operation for every ξ_i , we produce a sequence S^* , of the same length as S , composed of at most $4m$ different symbols.

The assertion of the theorem is then an immediate consequence of the following.

Lemma 4. S^* is a $(4m, s+2)$ -sequence.

Proof. Since it is clear that no two adjacent elements of S^* can be equal, it remains to show that S^* does not contain an alternating subsequence of the form $\mu \dots \nu \dots \mu \dots \nu \dots$ of length $s+4$. Suppose to the contrary that S^* does contain such an alternation, and consider any four consecutive elements of this alternation. Choose non-vertex points $x, y \in \mu$ and non-vertex points $z, w \in \nu$ so that C passes through these points in the order x, z, y, w . Consider the following Jordan arcs (see Figure 3).

β_{xy} is an arc inside $\overline{\mu}^*$ connecting x to y ;

β_{zw} is an arc inside $\overline{\nu}^*$ connecting z to w ;

C_{xy} is the portion of C traversed in direction from x to y with face f on the left.

Note that C_{xy} does not intersect β_{xy} and β_{zw} except at their endpoints. We claim that β_{xy} and β_{zw} must intersect one another.

The union of C_{xy} and β_{xy} forms a closed Jordan curve J which cuts the plane into two components. Call the one that contains the face f the outside. Since w does not lie on J and lies on the boundary of f , it must lie outside J . Since z lies in the relative interior of an arc that is common to J and the boundary of f there must be some z' on β_{zw} sufficiently close to z that does not lie outside J . Thus the portion of β_{zw} that forms a path between z' and w must intersect J . As it does not intersect C_{xy} , it must intersect β_{xy} , as claimed.

This shows that each quadruple of consecutive elements in our alternation induces at least one intersection point between the corresponding arcs $\beta_{xy} \subset \mu$ and $\beta_{zw} \subset \nu$. Moreover, it is easily checked that for any pair of distinct quadruples of this type, either the two corresponding subarcs of the form β_{xy} along μ are disjoint, or the two subarcs β_{zw} along ν are disjoint. Thus all these intersections must be distinct. Since the number of such quadruples is $s+4-3=s+1$, we obtain a contradiction, which completes the proof of the lemma. ■

Lemmas 3 and 4 complete the proof of Theorem 2, and, as argued above, also the proof of Theorem 1. As corollaries to Theorem 1, we obtain

- (a) $\sigma_n(\mathbf{S}) = \Theta(n\alpha(n))$ for \mathbf{S} the collection of all line segments.
- (b) $\sigma_n(\mathbf{C}) = O(n \cdot 2^{\alpha(n)})$ and $\sigma_n(\mathbf{C}) = \Omega(n\alpha(n))$ for \mathbf{C} the collection of all circles.
- (c) A similar upper bound holds for arrangements of the boundaries ∂K_i of n vector differences $K_i = A_i - B$, $i=1,2,\dots,n$, for disjoint closed convex sets A_i and a closed convex set B , as discussed in the introduction.

The proof of all these claims is straightforward. All upper bounds are immediate consequences of Theorem 1 and of the observations made in the introduction. The lower bound in (a) for line segments follows from the fact, proved in [WS] and [Sho], that even a single face in an arrangement of n segments can contain $\Omega(n\alpha(n))$ edges. The lower bound in (b) for circles follows from a recent result of [Sho] giving a construction of n circular arcs whose endpoints lie on the x -axis and which otherwise lie above the x -axis, such that their lower envelope consists of $\Omega(n\alpha(n))$ subarcs. If we complete each of these arcs to a full circle, and approximate the x -axis by a sufficiently large circle, we obtain the claimed lower bound. See also the remarks at the end of Section 3 below for other related results.

Another interesting consequence of Theorem 1 can be stated as follows. Take $\Gamma = \{\gamma_1, \gamma_2, \dots, \gamma_n\} \subset \Gamma_0$ and consider the zones defined by γ_i in $A(\Gamma - \{\gamma_i\})$ for all $1 \leq i \leq n$. We take the sum of the numbers of edges of all n zones which we know is $O(n\lambda_{s+2}(n))$, with s defined as usual. Let f be a face in $A(\Gamma)$, let k_f be the number of curves contributing edges to f , and let $|f|$ be the number of edges bounding f . Observe that the edges of f are counted for every one of the k_f curves. More precisely, when we consider the zone of γ_i which contributes j edges to f , then we add at least $|f| - c \cdot j$ to the total sum, where c is some positive constant between 1 and 2. Here we get a constant c not necessarily equal to 1 because two edges of f that are separated by an edge on γ_i can belong to a single edge in the face of $A(\Gamma - \{\gamma_i\})$ that contains f . Thus, the total contribution of f to the above sum is at least $k_f |f| - 2|f|$. Since the sum of the $2|f|$ taken over all faces f of the arrangement is $O(n^2)$ we get

$$\sum_{f \text{ in } A(\Gamma)} k_f |f| = O(n\lambda_{s+2}(n)).$$

As shown in Theorem 2, $|f| = O(\lambda_{s+2}(k_f))$. This together with $\frac{\lambda_{s+2}(k_f)}{k_f} \leq \frac{\lambda_{s+2}(n)}{n}$ implies

$$\sum_{f \text{ in } A(\Gamma)} |f|^2 = O\left(\sum_{f \text{ in } A(\Gamma)} |f| \lambda_{s+2}(k_f)\right) = O\left(\frac{\lambda_{s+2}(n)}{n} \sum_{f \text{ in } A(\Gamma)} |f| k_f\right) = O(\lambda_{s+2}^2(n)).$$

Suppose we are now interested in the maximum number of edges bounding some m faces of $A(\Gamma)$. Using standard inequalities (see e.g. a similar analysis in [Ed]), we get

$$\sum_{i=1}^m |f_i| = O(m^{1/2} \lambda_{s+2}(n)).$$

We formulate this result as a theorem. It is weaker than bounds obtained for special

cases such as lines, line segments, and circles (see [EGS] and [CEGSW]) but applies to more general curves.

Theorem 5. Let Γ be a set of n curves from Γ_0 (satisfying the conditions mentioned before Theorem 1). The maximum number of edges bounding m faces of $A(\Gamma)$ is

$$O(m^{1/2}\lambda_{s+2}(n)).$$

3. Incremental Construction of General Arrangements

Let Γ_0 be a family of Jordan curves or arcs as defined in the introduction. We assume that the curves in Γ_0 have relatively simple shape. In particular, we assume that each $\gamma \in \Gamma_0$ consists of at most p smooth portions and has at most q points of vertical tangency, for some small fixed integer constants p and q . Moreover, we assume that each of the following operations can be performed in time proportional to the size of the requested output.

- (i) Find the intersection points of two curves in Γ_0 .
- (ii) Given a vertical line segment, find its highest and lowest intersections with a curve $\gamma \in \Gamma_0$.
- (iii) Using some standard parametrization of curves in Γ_0 , determine, for any $\gamma \in \Gamma_0$ and any two points $x, y \in \gamma$, the relative order of x and y along γ in this parametrization.
- (iv) Given a curve $\gamma \in \Gamma_0$, find its points of vertical tangency.

Let $\Gamma = \{\gamma_1, \gamma_2, \dots, \gamma_n\}$ be an n -element subcollection of Γ_0 . We assume that any two curves in Γ intersect only transversally, in at most s points, and the arrangement $A(\Gamma)$ is simple. For each $m \leq n$ let $\Gamma_m = \{\gamma_1, \gamma_2, \dots, \gamma_m\}$. Our goal is to calculate $A(\Gamma) = A(\Gamma_n)$ incrementally, by starting with $A(\Gamma_1)$ (which has a trivial structure) and by adding the curves γ_i one at a time, obtaining progressively the arrangements $A(\Gamma_2), A(\Gamma_3), \dots, A(\Gamma_n)$. Each output arrangement A is assumed to be represented by the following data structure.

Each face of A is split into subcells by drawing a vertical segment from each vertex v of A (including endpoints and points of vertical tangency along the given arcs), and extending it until it meets the arcs lying directly above and directly below v . Each resulting subcell is like a trapezoid – it is bounded by two vertical segments (each of which may degenerate to a single point, a half-line, or a full line), and, at its top and bottom, by a portion of an arc. We maintain this collection of subcells as a refinement of A so that each cell contains pointers to its edges and each vertex (whether an original vertex of A or a new endpoint of one of the vertical segments) points to its incident edges. It will be important that each cell is bounded by only a constant number of edges. This can be maintained only if we abandon the idea of having the subdivision stored as a cell complex (where every edge is incident to

exactly two faces). An edge is thus directed (e.g. in the sense that it is associated with the cell that lies to its left) and can overlap with an arbitrary number of edges that lie on the same arc but are directed the other way. With each (directed) edge we store the cell it bounds and also the leftmost edge that overlaps it and is directed the other way.

Since each vertex of A induces at most two vertical segments, it is clear that the complexity of the refinement of A is proportional to the complexity of A itself. We refer to the refinement of A as its *vertical cell decomposition*.

The incremental calculation of the sequence of arrangements proceeds as follows. Suppose the arrangement $A_m = A(\Gamma_m)$ has already been constructed for some $m < n$ (and is represented by its vertical cell decomposition). Given the next curve $\gamma = \gamma_{m+1}$, we wish to calculate $A_{m+1} = A(\Gamma_{m+1}) = A(\Gamma_m \cup \{\gamma\})$ (again represented by its vertical cell decomposition). The first step is to locate the trapezoidal cell of A_m containing some initial point z_0 of γ . For this, assume that γ is an x -monotone arc; this can always be enforced by breaking γ into $O(1)$ pieces at its points of vertical tangency. Let $z_0 \in \gamma$ be its leftmost endpoint, if it exists. Otherwise choose $z_0 \in \gamma$ to the left of all its intersections with the arcs of Γ_m . By drawing the vertical line passing through z_0 and determining its two nearest intersections with the arcs in Γ_m , it is easy to determine the cell c_0 containing z_0 in $O(m)$ time.

Next we trace γ from z_0 to the right, and keep track of all the trapezoidal cells of A_m intersected by γ . This is done as follows. Since each trapezoidal cell has constant complexity, we can find in constant time the first intersection q of γ with the boundary of c_0 . If q lies on the upper or lower boundary of c_0 , then it is a new vertex of A_{m+1} . In this case we split c_0 into two subcells by drawing the vertical segment from q through c_0 , split the new cell c_1 into which γ enters by a similar segment, and continue the tracing of γ in c_1 . The only difficult step in this computation is to find c_1 . Remember that the vertical cell decomposition of A_m is not a cell complex which, for example, means that the upper and lower edges of c_0 can have an arbitrary number of bordering cells on the other side. Suppose γ intersects the upper edge of c_0 . In this case we have a pointer to the leftmost cell on the other side of the edge and we simply check whether this cell is c_1 ; if it is not we mark the cell and check the cell to its right (which is one of at most two cells and can thus be found in constant time). As we check and mark cells we let the upper edge of c_0 maintain a pointer to the leftmost yet unmarked cell on the other side in order to avoid going through the same list of cells again in case γ intersects the upper edge of c_0 more than twice.

Suppose next that q lies on the right vertical boundary β of c_0 , and let c_1 be the adjacent cell. Assuming general position, we can suppose that β passes through just one vertex v of A_m (which can also be an endpoint or a point of vertical tangency along some arc). For specificity, we assume that v lies above q . Then the portion of β below q is superfluous in A_{m+1} , since it corresponds to no vertex of that

arrangement. It needs to be deleted and therefore the lower portion of c_0 below γ must be merged with the lower portion of c_1 . We thus split c_0 into two subcells along γ , assign the truncated $\beta = vq$ as the right boundary of the upper subcell, but leave the lower subcell "open-ended" to the right, and record the fact that it needs to be merged with subsequent subcells. We continue the tracing of γ in this manner, keeping track of which side of γ (if any) contains an open-ended cell. When γ reaches a top or bottom cell boundary, both the open-ended cell and the cell on the other side of γ terminate (and are processed in the manner described above). When γ reaches another right cell boundary, either the open-ended cell terminates (so that it can now be assigned an appropriate right boundary), and the cell on the other side becomes open-ended, or the open-ended cell continues to be open-ended, to be merged with further subcells. Continuing to trace γ in this manner until reaching its right end-point or its final unbounded edge in A_{m+1} , we obtain the desired new arrangement A_{m+1} , properly represented by its vertical cell decomposition. See Figure 4 for an illustration of this procedure.

Let us now assess the total amount of time needed for updating arrangement A_m to obtain A_{m+1} in this manner. The work mainly consists of walking from one trapezoidal cell to the next and updating the adjacency structure of trapezoidal cells as new cells are being formed and old cells are merged. By our marking strategy we guarantee that the time is proportional to the number of trapezoidal cells into which the (original) cells of A_m that meet γ are decomposed. These cells are exactly those whose edges form the zone of γ and, by Theorem 1, the maximum number of edges of a zone is $O(\lambda_{s+2}(n))$. The number of trapezoids in the decomposition is proportional to the number of edges in the zone.

By what we said above it takes time $O(\lambda_{s+2}(m))$ to construct the vertical cell decomposition of A_{m+1} if the vertical cell decomposition of A_m is given. We thus obtain the main result of this section.

Theorem 6. The incremental procedure for calculating the arrangement $A(\Gamma)$ runs in $O(n\lambda_{s+2}(n))$ time and takes $O(n^2)$ storage.

As corollaries to Theorem 6, we obtain the following results.

- (a) One can calculate an arrangement of n line segments in time $O(n^2\alpha(n))$. Of course, one can construct the arrangement faster, namely in time $O(n\log n + k)$ where k is the number of intersecting pairs of line segments, using a different algorithm (see [CE]).
- (b) One can calculate an arrangement of n circles in time $O(n^2 \cdot 2^{\alpha(n)})$.
- (c) Similarly, one can calculate in time $O(n^2 \cdot 2^{\alpha(n)})$ the arrangement of the boundaries ∂K_i of n vector differences $K_i = A_i - B$, $i=1,2,\dots,n$, for disjoint closed convex sets A_i and a closed convex set B (all having simple shapes). Hence, within the same time bound, one can find an optimal placement of B which maximizes either the number

of objects A_i intersected by B , or more generally the sum of certain weights associated with these objects.

Remarks. (1) The superlinear lower bounds on the complexity of a zone in an arrangement of segments or of circles, as mentioned at the end of section 2, do not necessarily imply a superquadratic lower bound on the complexity of an incremental algorithm for the calculation of the arrangement. For example, taking the lower bound construction of [WS], [Sho], which gives a collection of n line segments whose lower envelope consists of $\Omega(n\alpha(n))$ subsegments, and then adding n additional long horizontal segments, each lying below this envelope and above the preceding segment, it is easily checked that the total number of edges traversed during our incremental algorithm is $\Omega(n^2\alpha(n))$. However, this bound arises only when the horizontal segments are inserted in increasing order of their height. Many other orders, such as a random order (see Mulmuley [M]), or the reverse of the above order, will perform much better in this case.

(2) McKenna and O'Rourke [MO] have independently obtained a special case of Theorems 1, 2 and 6 for arrangements of hyperbolas. Because of the special properties of their arrangements they were able to avoid many of the topological difficulties that we had to face.

(3) In certain special cases one can obtain slight improvements both in the combinatorial complexity of a zone and in the complexity of the above incremental algorithm. For example, in the case of unit-circles, the analysis given in [CL] shows that the portion of the zone of a unit-circle γ in an arrangement of n other unit-circles which lies *outside* γ is only $O(n)$. Using this fact, [CL] obtain an incremental algorithm in which only the "outer zone" of each newly added circle is being traversed, resulting in an overall $O(n^2)$ complexity. In contrast, when considering the entire zone, our techniques can be adapted to prove an upper bound of $O(n\alpha(n))$ on the complexity of a zone. Specifically, we split each unit-circle at its rightmost and leftmost points to obtain an upper semicircle and a lower semicircle. Any two upper semicircles intersect in at most one point, and similarly, any two lower semicircles intersect at most once. Thus, Theorems 1 and 2 imply that the complexity of the zone of a new unit-circle γ in the arrangement A^+ (A^-) of all upper (lower) semicircles is $O(\lambda_3(n))=O(n\alpha(n))$. Finally, we invoke the combination lemma of [GSS] to deduce that the zone of γ in the arrangement of A^+ and A^- superimposed on one another (which is to say in the arrangement of the original full unit-circles) is also $O(n\alpha(n))$. Combining this observation with the results of [CL] we see that if indeed the zone of a new unit-circle γ is superlinear in size, this can only be if its portion *within* γ is superlinear. We do not know whether this can really happen.

(4) A very similar situation arises in the case of the zone of an arbitrary closed convex curve γ in an arrangement of n lines. The complexity of the "outer zone" of γ is bounded by the maximum complexity of a single cell in an arrangement of $2n$ half-lines, which by a recent result of [ABP] is only $O(n)$. On the other hand, the

complexity of the “inner zone” of γ is bounded by the maximum complexity of the unbounded cell in an arrangement of n line segments in the plane which by Theorem 2 is $O(n\alpha(n))$. We conjecture that the inner zone of such a curve γ can indeed have $\Omega(n\alpha(n))$ complexity.

An interesting problem for further study is whether the new technique of topological sweeping, as given in [EG] for arrangements of line, can be adapted to apply to arrangements of more general arcs and curves.

4. The Special Case of Lines

In this section we conclude with a relatively easy “exercise” which gives a new proof of the zone theorem for lines, using (a simplified version of) the general technique developed in this paper.

Let $L = \{l_1, l_2, \dots, l_n\}$ be a collection of n lines in the plane, and let l be another line, which without loss of generality is assumed to be the x -axis. Consider the edges of the zone of l in $A(L)$ that lie above l . As in section 2, we truncate each line l_i to its portion ρ_i lying in that half-plane, which is simply a half-line emerging from some point on the x -axis. We also expand each ρ_i to a narrow angular wedge ρ_i^* from that point, and distinguish between the left side ρ_i^- and the right side ρ_i^+ of ρ_i^* . We need to bound the number of edges in the bottom unbounded face f of the arrangement $A(\{\rho_1^+, \dots, \rho_n^+, \rho_1^-, \dots, \rho_n^-\})$. As in section 3, we pick some connected component C of ∂f . It is easily checked that in this case C must be unbounded, and we traverse it so that f lies to our right (informally this is equivalent to a left-to-right traversal of the faces of $A(L)$ crossed by l). We now write down the sequence S of half-lines in the order they appear along C , but we split this sequence into two subsequences S^- and S^+ so that S^- (S^+) contains only the appearances of the left half-lines ρ_i^- (the right half-lines ρ_i^+).

We claim that S^- and S^+ are both $(n, 2)$ -sequences, and we will prove this for S^- . First note that S^- does not contain a pair of equal adjacent elements. Indeed, suppose to the contrary that S^- contains an adjacent pair $\rho_i^- \rho_i^-$. Then S had to contain some other right half-line ρ_j^+ between these two left elements, and then C must also traverse some left half-line (other than ρ_i^-) between these two portions of ρ_i^- , a contradiction.


Next we show that S^- does not contain any alternating quadruple of the form $\rho_i^- \dots \rho_j^- \dots \rho_i^- \dots \rho_j^-$. Indeed, suppose to the contrary that S^- does contain such a quadruple. Let $x, y \in \rho_i^- \cap C$, $z, w \in \rho_j^- \cap C$ be four points appearing along C in the order x, z, y, w . Using the same arguments as in the proof of Lemma 4, one can show that the segments xy and zw must intersect (at the unique point q of intersection of ρ_i^- and ρ_j^-). But then it is easily checked that the angular wedge xqw must be disjoint from f , and that w cannot appear along C (see Figure 5).

Thus the length of S^- and S^+ is at most $2n_C-1$ each, where n_C is the number of half-lines appearing in C . Hence the total number of edges in C is at most $4n_C-2$. Summing over all components C , it follows that the total size of the "upper zone" of l is at most $4n-2$, which matches the bounds obtained in [EOS] and [CGL].

References

- [ABP] P. Alevizos, J.-D. Boissonnat and F. P. Preparata, On the boundary of a union of rays, to appear in *Proc. 6th Ann. Sympos Theoret. Aspects of Comput. Sci.*, 1989.
- [ASS] P. Agarwal, M. Sharir and P. Shor, Sharp upper and lower bounds on the length of general Davenport-Schinzel sequences, to appear in *J. Combin. Theory, Ser. A*.
- [AS] B. Aronov and M. Sharir, Triangles in space or building (and analyzing) castles in the air, *Proc. 4th Ann. ACM Symp. on Computational Geometry*, 1988, pp. 381-391.
- [CE] B. Chazelle and H. Edelsbrunner An optimal algorithm for intersecting line segments in the plane, *Proc. 29th Ann. IEEE Sympos. Found. Comput. Sci.*, 1988, pp. 590-600.
- [CGL] B. Chazelle, L. Guibas and D.T. Lee, The power of geometric duality, *BIT* 25 (1985) pp. 78-90.
- [CL] B. Chazelle and D.T. Lee, On a circle placement problem, *Computing* 36 (1986), pp. 1-16.
- [CEGSW] K. Clarkson, H. Edelsbrunner, L. Guibas, M. Sharir and E. Welzl, Combinatorial complexity bounds for arrangements of curves and surfaces, *Proc. 29th Ann. IEEE Sympos. Found. Comput. Sci.*, 1988, pp. 568-579.
- [Ed] H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer-Verlag, Heidelberg, 1987.
- [EGS] H. Edelsbrunner, J. L. Guibas and M. Sharir, The complexity of many faces in arrangements of lines and of segments, *Proc. 4th Ann. ACM Symp. on Computational Geometry*, 1988, pp. 44-55.
- [EOS] H. Edelsbrunner, J. O'Rourke and R. Seidel, Constructing arrangements of lines and hyperplanes with applications, *SIAM J. Comput.* 15 (1986), pp. 341-363. 1983, pp. 83-91.
- [GSS] L. Guibas, M. Sharir and S. Sifrony, On the general motion planning problem with two degrees of freedom, *Proc. 4th Ann. ACM Symp. on Computational Geometry*, 1988, 289-298.
- [HS] S. Hart and M. Sharir, Nonlinearity of Davenport-Schinzel sequences and of generalized path compression schemes, *Combinatorica* 6 (1986), pp. 151-177.
- [HMRT] K. Hoffman, K. Mehlhorn, P. Rosenstiehl and R.E. Tarjan, Sorting Jordan sequences in linear time, using level-linked search trees, *Information and Control* 68 (1986) pp. 170-184.
- [KO] Y. Ke and J. O'Rourke, Moving a ladder in three dimensions: Upper and lower bounds, *Proc. 3rd Ann. ACM Symp. on Computational Geometry*, 1987, pp. 136-146.
- [KLPS] K. Kedem, R. Livne, J. Pach and M. Sharir, On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles, *Discrete Comput. Geom.* 1 (1986), pp. 59-71.
- [MO] M. McKenna and J. O'Rourke, Arrangements of lines in 3-space: A data structure with applications, *Proc. Ann. 4th ACM Symp. on Computational Geometry*, 1988, 371-380.
- [M] K. Mulmuley, A fast planar partition algorithm I, *Proc. 29th Ann. IEEE Sympos. Found. Comput. Sci.*, 1988, pp. 580-589.
- [PSS] R. Pollack, M. Sharir and S. Sifrony, Separating two simple polygons by a sequence of translations, *Discrete Comput. Geom.* 3 (1988), pp. 123-136.

- [SS1] J.T. Schwartz and M. Sharir, On the piano movers' problem: I. The case of a rigid polygonal body moving amidst polygonal barriers, *Comm. Pure Appl. Math.* 36 (1983), pp. 345-398.
- [SS2] J.T. Schwartz and M. Sharir, On the two-dimensional Davenport Schinzel problem, to appear in *J. Symbolic Comput.*
- [Sh1] M. Sharir, Almost linear upper bounds on the length of general Davenport-Schinzel sequences, *Combinatorica* 7 (1987) pp. 131-143.
- [Sh2] M. Sharir, Improved lower bounds on the length of Davenport Schinzel sequences, *Combinatorica* 8 (1988), pp. 117-124.
- [Sho] P. Shor, Geometric realizations of superlinear Davenport Schinzel sequences, in preparation.
- [WS] A. Wiernik and M. Sharir, Planar realization of nonlinear Davenport Schinzel sequences by segments, *Discrete Comput. Geom.* 3 (1988), pp. 15-47.



$$S = (\delta_1^+ \delta_2^- \delta_2^+ \delta_1^+ \delta_6^- \delta_7^- \delta_7^+ \delta_6^- \delta_5^- \delta_4^+ \delta_3^+ \delta_2^- \delta_1^-)$$

Figure 1.






Figure 2.



— curves of Γ_m
 - - - new curve γ

vertical boundaries of the
 cell decomposition of A_m

new vertical boundaries

old vertical boundaries that
 are deleted in the new
 decomposition, making the
 cells they bound "open-ended"

Figure 4.




Figure 3.




Figure 5.

~ ?

c.2

DATE DUE

BORROWER'S NAME

FOURTEEN DAYS

A fine will be charged for each day the book is kept overtime.

GAYLORD 142			PRINTED IN U.S.A.

